Campus de l'Université de Montréal
C.P. 6128, succursale Centre-ville
Montréal QC H3C 3J7 Canada

Telephone (514) 343 7575
Télécopieur (514) 343 7121

crt@crt.umontreal.ca
www.crt.umontreal.ca

HEC MONTRÉAL · POLYTECHNIQUE MONTRÉAL · Université de Montréal

CRT-2001-45

Centre de recherche sur les transports (C.R.T.)
Centre for Research on Transportation

EXPLOITING KNOWLEDGE ABOUT FUTURE DEMANDS FOR REAL-TIME VEHICLE DISPATCHING

by

Soumia Ichoua
Michel Gendreau
Jean-Yves Potvin

November 2001

103-03077/

PUBLICATION
CRT-2001-45

EXPLOITING KNOWLEDGE ABOUT FUTURE DEMANDS FOR REAL-TIME VEHICLE DISPATCHING*

by

Soumia Ichoua[1]
Michel Gendreau[1]
Jean-Yves Potvin[1]

* This research was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC), and by the Fonds pour la Formation de Chercheurs et l'Aide à la Recherche (FCAR).

[1] Département d'informatique et de recherche opérationnelle and Centre de recherche sur les transports, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, Canada H3C 3J7

Dépôt légal – Bibliothèque nationale du Québec
Bibliothèque nationale du Canada, 2001

ABSTRACT

An important, but seldom investigated, issue in the field of dynamic vehicle routing and dispatching is how to exploit information about future events to improve decision-making. In this paper, we address this issue in a real-time setting with a strategy based on probabilistic knowledge about demand arrivals to better manage the fleet of vehicles. More precisely, the new strategy introduces dummy customers (representing forecasted requests) in vehicle routes to insure a good coverage of the territory. This strategy is assessed through a set of computational experiments performed in a simulated environment.

Key words : vehicle dispatching, real-time, anticipatory, probabilistic knowledge, heuristic, tabu search, parallelism.

RÉSUMÉ

L'exploitation de la prévision des demandes futures demeure l'une des questions encore ouvertes dans le domaine de la répartition des véhicules en temps réel. Dans cet article, nous nous attardons sur ce sujet en proposant une stratégie qui se base sur une connaissance probabiliste des demandes futures afin d'orienter judicieusement les véhicules de manière à subvenir aux besoins futurs. En particulier, la stratégie proposée introduit des demandes fictives (représentant des requêtes prévues) dans les routes des véhicules afin d'assurer une bonne couverture du territoire. Une évaluation des bénéfices de l'inclusion d'une telle stratégie dans le processus de prise de décisions a été effectuée à l'aide de simulations qui s'inspirent de situations observées dans les systèmes de courrier rapide. Les résultats démontrent que la stratégie proposée apporte des améliorations substantielles au niveau de la qualité des solutions.

Mots clés : répartition de véhicules, temps réel, anticipation, connaissance probabiliste du futur, heuristique, recherche avec tabous, parallélisme.

1 Introduction

The rapid growth in communication and information technologies, as well as the increased interest in just-in-time distribution systems, have recently led researchers to focus on dynamic vehicle routing and scheduling problems. In the static version of these problems, the goal is to design a set of minimum cost routes to satisfy the demand while satisfying various constraints. In a real-time setting, a number of service requests arise continuously and randomly over time. Thus, each new request must be inserted in real-time in the current set of planned routes (where a planned route is a sequence of requests that have been assigned to a vehicle, but have not been serviced yet). Clearly, any solution is based on incomplete information, since some requests have not arisen yet. Furthermore, some of the request attributes, like actual demand or time to start service, may be prone to random variations over time. In this context, uncertainty is an inherent characteristic of the problem.

Fortunately, future events often arise in a somewhat predictable way (e.g., historical averages and probability distributions can be derived from previously collected data). An important issue that has been rarely investigated in the field of dynamic vehicle routing and dispatching is how to exploit probabilistic knowledge about future events to improve the decision-making process. Actually, exploiting information about future events is likely to result in a better management of the resources at the current time, which implies that future needs will be met more efficiently. In this paper, we propose a strategy to exploit probabilistic information about future demand arrivals. The effectiveness of this strategy is assessed within a parallel tabu search heuristic, previously reported in Gendreau et al. (99). This development is motivated from the local operations of long-distance express mail services (e.g., Federal Express). In this application, customers call a central office when they want some mail to be delivered abroad. Local collection routes are designed to address these customers' requests. The collected mail is then brought back to the central office for further processing and shipping.

The static version of this problem can be stated as follows. A fixed size fleet

of m identical vehicles is used to service customer requests. The goal is to design a set of minimum cost routes which service every customer, such that:

- each vehicle route originates from and terminates at a central depot;
- each vehicle services one route;
- each request is served once by exactly one vehicle;

Furthermore, each service point i (including the depot) has its own time window $[e_i, l_i]$, where e_i is the earliest service time and l_i is the latest service time. These time windows are "soft", i.e., a vehicle can arrive before the lower bound or after the upper bound. If the vehicle is too early, it must wait to start its service; if the vehicle is too late, a penalty for lateness is added to the route cost (which is usually related to the total distance traveled by the vehicle). However, the upper bound of the time window at the depot is hard. Thus, a vehicle is not allowed to come back late.

The paper is organized as follows. Section 2 provides a brief literature review dedicated to stochastic vehicle routing problems in static and dynamic settings. In this section, we mostly focus on the exploitation of knowledge about future demands. Section 3 introduces a new probabilistic strategy for managing the fleet of vehicles and discusses different issues related to it. Section 4 briefly describes the original tabu search heuristic reported in Gendreau et al. (99) and explains how it was modified to integrate the new strategy. Section 5 reports computational results obtained with the new algorithm. Finally, section 6 summarizes the major findings of the paper.

2 Literature review

In the field of vehicle routing, the issue of exploiting knowledge about future demands to improve decision-making has gained an increasing interest in the last decade. This issue has been addressed in the literature under two different problem classes (which are associated with two different ways of tackling the uncertainty, one static, the other dynamic). The existing work for each class is reported in the two following subsections.

2.1 Static stochastic vehicle routing problems

In this problem class, the customers and/or their demand are discrete or continuous random variables. Probabilistic information about future events is used to construct an "a priori" solution which optimizes the expected value of a given objective function. The probabilistic traveling salesman problem (PTSP), the vehicle routing problem with stochastic demands (VRPSD), the vehicle routing problem with stochastic customers (VRPSC), and the vehicle routing problem with stochastic customers and demands (VRPSCD) all belong to this class. A survey on these problems can be found in Gendreau et al.(96). In the following, these problems are briefly described and an overview of some related work is provided.

The PTSP is defined on a graph $G = (N, A)$, where N is the set of nodes (customers) and A is the set of arcs. A cost c_{ij} is associated with each arc $(i, j) \in A$ and each customer i has a probability p_i of requiring service. The goal is to find a single tour of minimum expected cost. This problem was first introduced by Jaillet (85, 88) and investigated by many researchers (e.g., Jezequel (85), Bertsimas and Howell (93) and Laporte et al. (94)). In the VRPSD, vehicles of fixed capacity must service customers with a stochastic demand. Some recent papers have been devoted to this problem (Bertsimas (88), Laporte et al. (89), Séguin (94), Gendreau et al. (95) and Yang et al. (2000)). In the VRPSC, each customer i has some probability p_i of requiring service, but the demand is deterministic (Bertsimas (88) and Waters (89)). The VRPSCD combines characteristics of the VRPSC and the VRPSD, given that both the customers and their demand are stochastic. This problem was investigated in Bertsimas et al. (90), Bertsimas (92), Gendreau et al. (94, 95) and Séguin (94).

In these problems, it is assumed that the probability distributions are known and are used to construct a solution in two stages. In a first stage, a planned (or "a priori") route is constructed which minimizes an objective function related to some expected measure (e.g., expected length of the routes). When the routes are actually executed in a second stage, recourse actions are applied to the first stage solution when needed to respond to the current situation (e.g., "jumping over"

a customer who does not show up). The interested reader will find more details about this kind of solution approach in Dror et al. (89) and Gendreau et al. (96).

2.2 Dynamic stochastic vehicle routing problems

Here, new events are handled as they are revealed over time. Thus, no solution is constructed "a priori" in this case. Assuming that some probabilistic knowledge about future events is available, two different ways of exploiting this information are reported in the literature. This is the subject of the following subsections.

2.2.1 Analytical studies

Here, the only source of uncertainty considered is the occurrence of new requests and/or the service (or dwell) times at customers. The probability distribution of these random variables is used to perform analytical studies aimed at examining the robustness and asymptotic behavior of different deterministic and myopic dispatching strategies. Bertsimas and Simchi-Levi (96) pointed out that this kind of work also provides new insights into the solution structure, thereby offering valuable guidelines for the design of new efficient deterministic algorithms. An example is found in Bertsimas and Van Ryzin (91), where the authors consider the dynamic traveling repairman problem (DTRP). In this problem, a vehicle is used to service a set of independently and uniformly distributed customer demands which occur according to a Poisson process with intensity parameter λ. The service time at each customer is also a random variable. The goal is to find a routing policy which minimizes the expected system time, which is the time spent in the system by each customer. Using results from geometrical probability, combinatorial optimization and queuing theory, the authors provide an analytical study of the problem in the case of light traffic and heavy traffic. First, the authors derive lower bounds for the average system time. Then they analyze several simple policies (e.g., First Come First Served) and compare their performance with these lower bounds. In 1993, the same authors (Bertsimas and Van Ryzin (93)) extend their findings to the case of a homogeneous fleet of m vehicles where each vehicle can visit at most q customers before returning to the depot. The uncapacitated version of this prob-

lem was also investigated. The authors partition the given area into subregions and use the results of the first paper to propose and analyze new policies in case of light and heavy traffic, respectively.

Swihart and Papastavrou (99) introduce a new variant of the DTRP where each service request has a pick-up and a delivery location. The objective is to minimize the expected system time. The authors consider the unit-capacity case where the vehicle can carry no more than one item, as well as the case where the vehicle can carry an arbitrary large number of items. In each of these two cases, lower bounds are derived for the average system time in light and heavy traffic conditions. Several policies are also proposed and compared.

In his dissertation, Rubio Maqueo (1995) has studied the vehicle routing and inventory problem (VRIP). In this problem, a single capacitated vehicle operates out of a single warehouse and services a finite set of retailers. The arrival of demands at a given retailer is modeled as an independent renewal process while the quantity of demand is stochastic. Also, the travel times between retailers are random variables. The objective is to minimize the expected inventory and transportation costs. Assuming that the system operates in heavy traffic, the author analyzes two models. In the first one, the vehicle follows a predefined routing scheme: deliver full loads to either a single retailer (direct shipping (DS)), or along a pre-specified TSP tour. In the second model, which is based on the current inventory levels, the controller dynamically chooses whether the vehicle should follow a TSP or a DS routing scheme. Using simulations and queuing theory, the author fully characterizes a dynamic control policy which is asymptotically optimal for the fixed routing scheme. In addition, the author investigates the importance of various operational decisions (e.g., TSP vs. DS, fixed vs. dynamic routing), thereby giving insights into the nature of the optimal solution to the VRIP.

In the first part of his dissertation, Kleywegt (96) introduces and studies a problem called the dynamic and stochastic knapsack problem (DSKP). In this problem, demands for a given resource occur according to some stochastic process. Each demand requires a random amount of resource and has a random profit. Demands can be either accepted or rejected. If a demand is accepted, an associated reward

is gained. Conversely, if it is rejected, a penalty is incurred. There is also a deadline, possibly infinite, after which demands can no longer be accepted (although the process can be stopped at any time before this deadline). The objective is to determine an optimal policy for accepting the demands and for stopping the process while maximizing the expected reward value. The author considers both the infinite and finite horizon cases. He establishes a number of theoretical characteristics for the optimal accepting rule and for the optimal value function, and uses these characteristics to derive an optimal policy for the problem. In the second part of his thesis, the author formulates and analyzes the dynamic and stochastic distribution problem (DSDP) with a network of terminals and a fleet of m vehicles. In this problem, customers call at random times for the transportation of a set of loads from given origins to given destinations. Each accepted request generates a reward and each rejected request incurs a penalty. There is also a holding cost per unit of time which depends on the number of vehicles, as well as the number of loads waiting at each terminal. The objective is to maximize the expected total value (rewards minus costs). The problem is formulated as a Markov decision process. Equations are derived for the optimal expected value and some fundamental structural characteristics of the solution are established. A two-stage algorithm is also proposed. The first stage is a threshold acceptance rule which is shown to be optimal, and the second stage is an optimal dispatching rule.

2.2.2 Stochastic approaches

This line of research has sought to develop explicit stochastic and dynamic approaches which incorporate the current information and probabilities of future events to produce more efficient solutions than those obtained through myopic deterministic strategies. The only work that we are aware of in this category is that of Powell and his team (Powell (88), Powell (96), Powell et al. (2000)), and a recent research done by Larsen (2000). In the following, we briefly discuss these studies.

The problem considered by Powell and his team is motivated from long-haul truckload trucking applications. In this context, the problem consists in dynamically

assigning drivers to loads which arise randomly over time. Each load is characterized by its origin, its destination and its time window. At any time, a vehicle can be either empty or carrying a single load. Furthermore, a demand may be rejected if it cannot be serviced within reasonable time. Powell (88) addresses the issue of relocating vehicles in anticipation of future demands. He first proposes a dynamic network model based on a time-space graph where the nodes correspond to known and forecasted demands in different zones at different time periods. In addition to the arcs which represent moves between regions, a dummy arc from each origin to a super-sink captures the impact of the current time period on the future. The cost on this arc is the expected reward value of having an additional vehicle in the associated region at the beginning of the next time period. This model has some limitations since a forecasted load may be serviced even if it never materializes. To overcome this weakness, the author proposes a stochastic, dynamic model where the demand $D_{ij}(t)$ between nodes i and j at time period t is now a random variable. Hence, the flow $x_{ij}(t)$ of loaded vehicles as well as the flow $y_{ij}(t)$ of empty vehicles moving between i and j at the current time period are random variables. This model is written as a two-stage stochastic program where the objective function combines a deterministic cost that corresponds to the first time period (the present) and the value of the recourse function that represents the expected cost from the next time period to the end of the time horizon (the future).

In Powell (96), the author investigates in the same context a hybrid model for solving a combined assignment and fleet management problem. The model is based on a dynamic-stochastic network with two components. The first one is the assignment network aimed at assigning drivers to loads during the current time period. The second is the forecast network which represents the remainder of the time horizon and includes all forecasted demands, as well as known loads to be picked up in the future. Arcs in this part of the network represent loaded and empty moves or a waiting time from one period to the next. A set of links connecting the two parts of the network represent empty relocation moves in anticipation of future loads. The expected recourse function is represented by a cluster of "recourse

links" which capture the expected marginal contribution of available vehicles in a given region during a given time period. An approximation of this function value is added to the arc costs and the resulting transportation problem is then solved with a network simplex algorithm on a rolling horizon basis.

In Powell et al. (2000), a new approach is proposed for solving large-scale stochastic and dynamic problems such as integer multi-commodity flow problems. In these problems, new decisions are made continuously over time in response to new information. The approach is first presented within a general framework. Then, the authors report its application in different contexts, like vehicle dispatching. The proposed approach combines techniques from dynamic programming and multistage stochastic programming. More precisely, it uses forward dynamic programming techniques (Bertsekas and Tsitsiklis (96)) that step forward in time to avoid the evaluation of all possible states. The curse of dimensionality, which is inherent to dynamic programming approaches, is thus alleviated.

The issue of relocating an empty vehicle in anticipation of future demands is also addressed in Larsen (2000), where probabilistic knowledge about future events is exploited. The problem studied is a dynamic traveling salesman problem where the goal is to find a minimum cost tour through a set of dynamic requests. The requests, each with a soft time window, unfold over time in n subregions according to a Poisson process with rate $\lambda_i, i = 1, ..., n$. Each subregion may have one or more resting locations (idle points) for the vehicles. Moreover, all requests must be serviced even if some vehicles arrive late at the depot. The objective function is a weighted sum of total travel time and total lateness over all requests. Whenever a new service request occurs, the current tour is re-optimized using a simple heuristic based on an exchange procedure. Once the vehicle has finished serving its current customer, the algorithm determines the best resting location for the vehicle, if the time window at its next destination does not allow immediate service (and assuming that the probability of having at least one request in the chosen location is high enough). This is achieved using the following on-line policies: **(i)** the vehicle goes to the nearest resting location, **(ii)** the vehicle goes to the busiest resting location (i.e., in the subregion with the largest λ_i value), **(iii)** the vehicle goes

to the resting location which belongs to the subregion with the highest expected number of customers. Simulations on problems with about 40 customers were used to compare the different relocation policies including a "no-rest" policy where the vehicle systematically waits at its current position. Unfortunately, the empirical analysis does not clearly identify a dominant policy.

3 A probabilistic strategy for managing the fleet of vehicles

3.1 Motivation

The challenge of designing models and methods which accurately reflect real-life conditions is an on-going strive in the field of dynamic vehicle dispatching. An important concern in this line of study is the exploitation of knowledge about future events to improve solution quality. In real-life, skilled dispatchers typically end up with some valuable knowledge about market demand (e.g., "peak" time periods, intense geographic zones, etc...). This knowledge allows them to better manage the current resources in anticipation of forecasted needs. Hence, an interesting issue in the field of dynamic vehicle dispatching, is the development of "well informed" solution approaches which explicitly reproduce the dispatchers' behavior. This has been investigated by Powell (86, 96) in the context of a truckload carrier (see section 2). In this case, the problem is a combined pick up and delivery problem with no consolidation (i.e., at any time, a vehicle is either empty or is carrying a single load) and the horizon typically extends over a number of days.

With regard to the work of Larsen (2000), we recall that when the vehicle finishes the service at its current customer, the two proposed policies choose, among all resting locations (including the current position), where the vehicle should spend its waiting time until the beginning of service at its next customer. In fact, it might be better for a vehicle to systematically wait at its current position rather than anywhere else. In the worst case (i.e., when the forecasted demands do not materialize), waiting at a resting location which is different from the current vehicle

location simply incurs an additional traveling distance.

As mentioned by Larsen, the proposed policies can also suggest a bad choice, namely to move the vehicle in a zone which is surrounded by many other zones with low request arrival rates, and far from regions where new demands are very likely to appear. This is illustrated in figure 1. In this figure, the number in each

2	1	4
3	1	1
3	3	2

Figure 1: Moving the vehicle far from an "interesting" area

square represents the arrival rate λ in the corresponding zone. In this case, the two policies would choose the zone with $\lambda = 4$, thereby ignoring the "interesting area" in the bottom. To overcome this shortcoming, one has to take into account the arrival rate in neighboring zones. In the following, we propose a probabilistic approach which addresses this issue in a broader context.

3.2 The strategy

The proposed strategy exploits probabilistic knowledge about the arrival of future demands to manage the fleet of vehicles more efficiently. The intent is to provide a good coverage of the territory in anticipation of forecasted demands.

3.2.1 Basic assumptions

Let us assume that the distribution area Z is partitioned into J zones $Z_1, Z_2, ..., Z_J$ and that the time horizon is divided into K time intervals $I_1, I_2, ..., I_K$. Moreover, we suppose that it is possible to estimate the probability p_{jk} that at least one demand will occur in zone Z_j during time period I_k. If this probability is high, it might be worthwhile to force the vehicle to wait in this zone instead of directing it

towards its next planned destination (if the latter is in another zone). The intent is to keep the vehicle readily available to service demands that are likely to unfold in its current zone.

To avoid directing the vehicle towards "bad" areas (as illustrated in figure 1), the probability that at least one demand will occur is calculated using a location based analysis. That is, instead of applying directly the probability p_{jk} of having at least one demand in the current zone Z_j, the probability is calculated over a neighborhood of the current vehicle position, which may have a non empty intersection with many zones. In the following we discuss the various elements which are taken into account by our new decision rule.

3.2.2 Waiting time

Actually, it can be useless for the vehicle to wait at its current location if the probability of occurrence of a new request in its neighborhood in the near future is not high enough. Moreover, while the vehicle is forced to wait, some new "real" requests in other zones, which appear further along its planned route, will suffer from service delays. To overcome this problem, a vehicle i is kept at its current position only for some amount of time δt_i. Parameter δt_i is very important: if it is too small, there may not be enough time for forecasted demands to materialize. If it is too large, the vehicle may wait for a long time for forecasted demands that will never come true, which may lead to additional lateness on its planned route.

In our implementation, $\delta t_i = min(\Delta, \bar{t}_i - t, (\underline{t}_i - t) + \eta)$. The first component Δ is a fixed parameter which is simply aimed at avoiding too large values. The second component is the difference between the latest feasible departure time $\bar{t}_i$ (which allows the vehicle to return before the upper bound of the time window at the depot) and the current time t. The third component is the difference between the latest departure time which does not induce any additional lateness on the vehicle route $\underline{t}_i$ and the current time t, plus a given parameter η. This last parameter reflects some tolerance to additional lateness.

3.2.3 Probability of occurrence of a new request

In our strategy, a given vehicle i may be forced to wait at its current position for some amount of time δt_i. But, because of the probabilistic nature of the available information, these demands may never materialize. In this case, the vehicle will wait for nothing. Thus, this waiting time is allowed only if the probability for a demand to occur in the vehicle neighborhood during time interval $[t, t+\delta t_i]$, where t is the current time, is greater than or equal to a given threshold. In the following, the probability calculation is explained.

We suppose that the service requests occur within a time-space Poisson process where λ_{jk} is the demand rate in zone Z_j during time period I_k, $1 \leq j \leq J$, $1 \leq k \leq K$. The Poisson distribution has some nice properties that lead to easily derived analytical expressions for describing the system. Let us assume that vehicle i has just finished servicing its current customer in zone Z_j. Then, the neighborhood $\bar{N}_i$ of vehicle i is defined as a square centered at the current vehicle position l_i. In our experiments, $\bar{N}_i$ has the same dimension as a zone, although it does not need to be in general. This is illustrated in figure 2 for a distribution area divided into 9 zones. In this figure, the current position l_i is represented by a large dot.

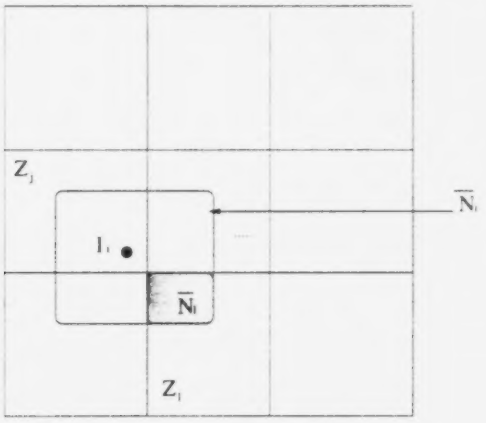

Figure 2: Defining the vehicle neighborhood $\bar{N}_i$.

Let also assume that at instant t, vehicle i is ready to leave its current location l_i

in zone Z_j. Then, the probability $p^{\delta t_i}$ that at least one demand will occur during time interval $[t, t+\delta t_i]$ in the neighborhood $\bar{N}_i$ is $p^{\delta t_i} = 1 - p_0^{\delta t_i}$, where $p_0^{\delta t_i}$ is the probability that no new request occurs during time interval $[t, t+\delta t_i]$ in $\bar{N}_i$. If we suppose that t belongs to time interval I_{k_1} and that $(t+\delta t_i)$ belongs to time interval I_{k_2} $(k_1, k_2 \in \{1, .., K\})$, then the interval $[t, t+\delta t_i]$ can be partitioned into $(k_2 - k_1 + 1)$ intervals, each of length δt_i^r and belonging to time period I_r $(k_1 \leq r \leq k_2)$. This is illustrated in figure 3 where the time horizon is partitioned into three time periods. In this figure, t belongs to time period I_1 and $(t+\delta t_i)$ belongs to time period I_3.

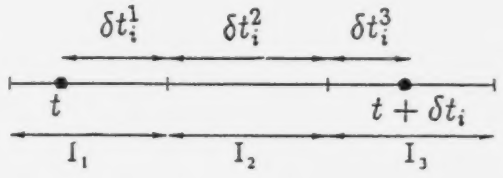

Figure 3: Partitioning δt_i.

The probability $p_0^{\delta t_i} = \prod_{r=k_1}^{k_2} p_0^{\delta t_i^r}$, where $p_0^{\delta t_i^r} = exp(-\lambda_{\bar{N}_i r} \delta t_i^r)$ is the probability that no new request occurs during δt_i^r in $\bar{N}_i$ and where $\lambda_{\bar{N}_i r}$ is the demand rate in $\bar{N}_i$ during time period I_r. Since the J Poisson processes are assumed to be independent, $\lambda_{\bar{N}_i r} = \sum_{l=1}^{J} \lambda_{lr} \frac{S_{\bar{N}_l}}{S}$.
In this equation, λ_{lr} is the demand rate in zone Z_l during time period I_r;
$\bar{N}_l = Z_l \cap \bar{N}_i$, $S_{\bar{N}_l}$ is the area of $\bar{N}_l$, and S is the area of a zone.

3.2.4 Coverage of the territory

While applying the proposed strategy, it may well happen that most or even all vehicles will be kept in a zone where the demand rate is high. Thereby, a new request which appears in another zone may be serviced very late or not serviced at all. To overcome this shortcoming, the number of vehicles in each zone must always be less than or equal to a given upper bound.

3.2.5 Summary of the strategy

Figure 4 gives a pseudo-code description of the proposed strategy for a given vehicle.

1. if $(d \leq \alpha\ \bar{d})$ go to the next customer location (and exit the procedure).
2. if $((prob[\text{at least one demand in neighborhood during } \delta t] > s)\ \&\ (v \leq V))$
   wait at the current location for δt time units,
   else go to the next customer location.

Figure 4: Pseudo-code of the proposed strategy.

In this figure, $\bar{d}$ is the average distance between two consecutive customers locations and d is the distance between the current position of the vehicle and the location of its next destination. Also, v is the number of vehicles in the current zone and V is the upper bound on the number of vehicles in this zone. Finally, α and s are user-defined parameters. Thus, a vehicle waits at its current location if its next destination is far enough, the probability of occurrence of a new request in the vehicle's neighborhood is high enough and there are not too much vehicles in the current zone.

The effectiveness of this rule is assessed by incorporating it within a parallel tabu search heuristic, previously reported in Gendreau et al. (99). The next section will briefly describe the original algorithm and explain how it was modified to integrate the proposed strategy.

4 A parallel tabu search algorithm

4.1 The original algorithm

Given the dynamic context, a solution at any instant t is a set of planned routes, each beginning with the current destination of the associated vehicle. This current destination is fixed in the procedure to be described and cannot be modified (i.e., the vehicle must reach it). The algorithm proposed in Gendreau et al. (99) is a parallel tabu search heuristic with adaptive memory (Rochat and Taillard

(95)). Tabu search is an iterative local search technique where a neighborhood is generated around the current solution and the best solution in this neighborhood becomes the new current solution (even if it does not provide any improvement). By allowing a degradation of the objective, it is possible to escape from bad local optima, as opposed to pure descent methods. The interested reader will find more details about this approach in Glover and Laguna (97).

The algorithm used in Gendreau et al.(99) can be summarized as follows :

- Construct I different initial solutions with a stochastic insertion heuristic (where the choice of the next customer to be inserted is randomized).
- Apply tabu search to each solution and store the resulting routes in the adaptive memory.
- While a stopping criterion is not met do :
  - Use the routes stored in the adaptive memory to construct an initial solution.
  - Decompose the problem into subproblems obtained through a geographical decomposition procedure.
  - Apply tabu search to each subproblem.
  - Add the routes of the resulting solution in the adaptive memory.

The objective is to minimize a weighted sum of total distance traveled and total lateness over all customers.

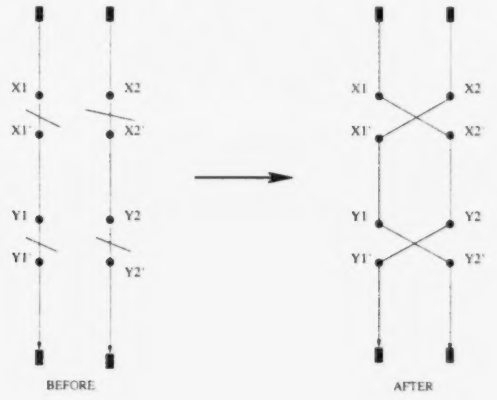

Figure 5: A CROSS exchange.

The procedure for generating the neighborhood of the current solution is called the CROSS exchange (Taillard et al. (97)). Basically, two segments of variable lengths are taken from two different routes and are swapped (see figure 5).
To speed up the algorithm, a parallel implementation on a network of workstations was developed. The parallelization of the procedure was achieved at two levels :
(1) Different tabu search threads run in parallel, each of them starting from a different initial solution.
(2) Within each search thread, many tabu searches run independently on the subproblems obtained through the decomposition procedure.
Whenever a new event occurs, the tabu search threads are interrupted. Then, after an appropriate update of the adaptive memory to reflect the current state of the system, the search is restarted with new solutions constructed from the updated memory. Thus, the tabu search is used to improve the current (static) solution between the occurrence of new events, where a new event may be either the occurrence of a new service request or the completion of service at a given customer (since a next destination must be indicated to the driver).

During the schedule generation, a least commitment strategy is considered. That is, if there is some time flexibility with regard to the lower bound of the time window at the next destination, the vehicle waits at its current location (rather than at its next destination).

4.2 Implementing the new strategy

The new strategy described in section 3.2.5 basically modifies the way of handling new events. This is explained in the following.
(a) When vehicle i has finished serving a customer, the new strategy is used to decide whether to wait or not. In the latter case, the vehicle is directed toward its next destination as in the original algorithm. In the former case, a dummy customer $D_{\delta t_i}$ is created at the current vehicle position. The intent is to force the vehicle to wait at its current location during time interval $[t, t + \delta t_i]$, where t is the current time and δt_i is calculated as in section 3.2.2, using the best solution in adaptive memory. Then, the dummy customer $D_{\delta t_i}$ is inserted at the beginning

of the planned route of vehicle i in every solution in the adaptive memory and solutions that become infeasible are discarded (because of the hard time window at the depot, it may not be possible to wait for δt_i time units at the current vehicle position). Finally, the search threads are restarted with new solutions constructed from the updated memory and the optimization is performed with the assumption that vehicle i will restart at time $t + \delta t_i$.

(b) When a new request occurs, we first look if it is found in the neighborhood of some waiting vehicles. Among all such vehicles, only those which can start service as soon as they reach the new request are kept. Then, we choose vehicle i for which the least additional cost is incurred in its current route. The new request then replaces the dummy customer $D_{\delta t_i}$. If no vehicle qualifies for the new request, the latter is handled as in the original algorithm. Note that the new strategy offers additional insertion points for a new incoming request. These are located between the current vehicle location and its next destination, for each waiting vehicle.

If vehicle i was not used to service any incoming request when time $(t + \delta t_i)$ is reached (either because no new request occurred in its neighborhood or because it was never chosen), the dummy customer $D_{\delta t_i}$ is removed from its route and the vehicle is sent to its next planned destination, as indicated by the best solution in the adaptive memory.

5 Computational results

5.1 Simulation framework

Simulations were performed to assess the effectiveness of the proposed strategy. Hence, a simulator was developed to produce different operating scenarios which reflect as closely as possible what is observed in the real world. In the following, we describe the components of this simulator.

5.1.1 Time horizon

The simulator first sets up a time horizon which represents the entire labor day. For example, a labor day running from 9h00 to 16h00 is represented by a time horizon of 300 minutes. Within this horizon, the simulator produces discrete time events which correspond to either a request arrival or the end of service at a given customer location. The new request arrivals take place only in the first 240 minutes of the time horizon, to allow vehicles to service their planned routes before returning to the depot.

5.1.2 Generation of new service requests

Typically, the arrival rate of new requests varies over time (e.g., peak hours) and space (e.g., densely populated areas). Hence, a time-space Poisson process is used. We recall that the distribution area is partitioned into J zones $Z_1, Z_2, ..., Z_J$ and that the time horizon is divided into K time intervals $I_1, I_2, ..., I_K$. Moreover, λ_{jk} denotes the demand rate in zone Z_j during time period I_k, $1 \leq j \leq J$, $1 \leq k \leq K$. The characteristics of a new service request are generated as follows. Given the current time interval I_k, the demand rate $\lambda_k = \sum_{j=1}^{J} \lambda_{jk}$ of the global distribution area Z (which is made of zones $Z_j, j = 1, 2, ..., J$) over time interval I_k is used to determine the arrival time of the next request. Then, the zone Z_j where the new request will appear is determined using the following cumulative probabilities $(P_j^k)_{0 \leq j \leq J}$:

$$P_j^k = \begin{cases} 0, & \text{if } j = 0 \\ \sum_{l=0}^{j} p_{lk} & \text{otherwise} \end{cases}$$

where $p_{lk} = \frac{\lambda_{lk}}{\sum_{j=1}^{J} \lambda_{jk}}$.

The selected zone Z_j is such that $P_{j-1}^k \leq u < P_j^k$, where u is drawn from a uniform distribution $U[0, 1]$. To generate the time window associated with the new request r, the earliest pick-up time e_r is first generated in the interval $[t, t + \Delta T_1]$, where t is the current time and ΔT_1 is a constant. Then, e_r is possibly adjusted to avoid

exceeding $\tilde{t} = l_0 - tr(r,0) - s_r$, where l_0 is the end of the day, $tr(r,0)$ is the travel time from r to the depot and s_r is the service or dwell time. Clearly, it is not possible to return to the depot on time if the service at request r starts after $\tilde{t}$. Once the earliest pick-up time e_r is determined, the latest pick-up time l_r is set to $e_r + \Delta T_2$, where ΔT_2 is a constant. The value l_r is adjusted similarly to e_r if necessary.

5.2 Testing scenarios

Different sets of problems were generated to perform numerical tests. In these problems, the fleet size is set to 4 or 6 vehicles with all vehicles moving at a constant average speed of 30 kms/hour. The service time is equal to 5 minutes at each customer. Parameters ΔT_1 and ΔT_2 are both set to 30 minutes. The time horizon is set to 300 minutes (i.e., 5 hours) and is divided into three time periods: morning, lunch time and afternoon. The lunch time period is one hour long and the two other ones are two hours long. The distribution area is a 5 km X 5 km square area and the depot is located at (3.3 km, 4.5 km). Nine "active" zones are considered in the distribution area and are partitioned into two categories: category 1 contains the intense central zone where the demand rate is high, while category 2 contains the remaining peripheral zones where the demand rate is low. Moreover, in each problem the set of requests is divided into two subsets. The first subset contains requests that are known at the beginning of the day (i.e., 25% of the total number of requests). The second subset contains requests that unfold in real-time. In this case, the arrival time of each request is calculated based on a 3x2 time-dependent demand rate matrix $(\lambda_{cT})_{1 \leq c \leq 2, 1 \leq T \leq 3}$, where λ_{cT} is the demand rate in a zone of category c during time period T. The entries are adjusted to create two different types of scenarios: scenario 1 where the total number of requests is about 180, and scenario 2 where the total number of requests is about 240. In the latter case, a higher degree of dynamism is obtained because there are more real-time requests over the same time horizon (in average, we obtain 27 requests per hour for scenario 1, and 36 requests per hour for scenario 2). The demand rate matrices for the two scenarios, in number of requests per hour, are

reported in tables 1 and 2, respectively. It is worth noting that in real-life, new requests rarely unfold in the intense zone (associated with downtown) during the lunch time period. Thus, the corresponding demand rate is set to 0 for the two scenarios during that time period. Tests are performed on 10 instances for each scenario.

	T_1	T_2	T_3
Intense zone	16.50	0.000	33.00
peripheral zones	2.04	2.04	2.04

Table 1: *Demand rate matrix in scenario 1.*

	T_1	T_2	T_3
Intense zone	19.98	0.000	39.96
peripheral zones	2.52	4.98	2.52

Table 2: *Demand rate matrix in scenario 2.*

5.3 Experiments

The experiments reported in this section were performed on a network of 3 SUN UltraSparc-IIi workstations (300 Hz). Each process is programmed in C++ and communication between the processes is handled by the Parallel Virtual Machine Software. As mentioned previously, the objective is to minimize the sum of total travel time and total lateness over all customers. The value of the parameters in the original tabu search algorithm are kept as in Gendreau et al. (99). Parameters related to our decision rule are examined in the following.

5.3.1 Preliminary tests

To determine the best values for parameters α, s, η, V and Δ used in our waiting strategy (see figure 4), tests were performed with 4 vehicles over a sample of 4 problems. Several values had to be tested for every problem in the sample. Thereby, a significant amount of computation time is required if the size of the sample is too large. On the other hand, this size has to be large enough to get

a correct calibration. Since scenarios 1 and 2 represent different degrees of dynamism, the tests were conducted separately for the two scenarios. Table 3 and 4 present a sample of results obtained with different combinations of values using scenario 1 and scenario 2, respectively. We show only the best combinations of values in these tables (i.e., less than 5 unserved customers on average for scenario 1 and less than 23 unserved customers on average for scenario 2). In these tests, parameter V was set to 1 for the peripheral zones, while different values were tested for the intense central zone. The four numbers in each entry in tables 3 and 4 are the average number of unserved customers, total travel time, total lateness and objective value, respectively. Based on the results obtained, the values $\alpha_1 = 0.20$, $s_1 = 0.80$, $\eta_1 = 200$ *sec.*, $\Delta_1 = 300$ *sec.* and $V_1 = 1$ for scenario 1 and $\alpha_2 = 0.40$, $s_2 = 0.80$, $\eta_2 = 100$ *sec.*, $\Delta_2 = 180$ *sec.* and $V_2 = 1$ for scenario 2 were finally kept for further investigations. In these experiments, we also found that δt_i is about 4.50 minutes for scenario 1 and 2.90 minutes for scenario 2.

5.3.2 Numerical results

The final tests were conducted on 10 instances for each scenario with a fleet size of 4 or 6 vehicles.

Scenario 1

Tables 5 and 7 report results obtained with the original algorithm (Gendreau et al. (99)), with 4 and 6 vehicles, respectively, while tables 6 and 8 report the results obtained with the new algorithm. The results show that the new algorithm performs well against the original one. More precisely:

Experiments with 4 vehicles

The new strategy leads to an improvement over 8 of the 10 problems tested. In the case of the two other problems, the results obtained with the new algorithm are as good as those obtained with the original one. Note that among the eight instances where an improvement was observed, the original algorithm succeeds in servicing all customers in five cases. In these cases the new strategy also reduces the objective value by 1.1 % to 12.2 %. In the three other instances, the number of unserved customers is reduced. In two cases, this reduction is achieved a the cost

of a marginal increase in the objective value. Note also that the objective value is improved by 4.1 % on average over the 10 problems.

Experiments with 6 *vehicles*

Since these problems are easier than those with only 4 vehicles, the original algorithm succeeds in servicing all customers in 9 problems among the 10 tested. On two problems, the results obtained with the new algorithm and the original one are the same. On the other hand, the new strategy improves the objective value on six problems. The reduction ranges from 1.1 % to 8.0 %. Moreover, when the original algorithm does not succeed in servicing all customers, the new strategy does. Note that the average improvement in the objective is 2.3 % over the 10 problem instances.

Scenario 2

The problems in this scenario are harder than those of scenario 1 due to the higher demand rates. Hence, the original algorithm never succeeds in servicing all customers. Tables 9 and 11 report results obtained with the original algorithm, with 4 and 6 vehicles, respectively, while tables 10 and 12 report the results obtained with the new algorithm. As expected, the new strategy clearly improves the results obtained with the original algorithm.

Experiments with 4 *vehicles*

In three problems, the two algorithms lead to the same number of unserved customers. However, the new strategy reduces the objective value from 3.0 to 7.5%. In two problems, the increase in the objective value is marginal when compared with the reduction in the number of unserved customers (c.f. 0.2 % against 3.0 % and 0.4 % against 6.3 %, respectively). On the other hand, the new algorithm succeeds in reducing both the number of unserved customers and the objective value for five problems. The reduction ranges from 2.0 % to 16.0 % for the objective value, and from 3.4 % to 29.0 % for the number of unserved customers. In these experiments, the average reduction is 4.2 % for the objective value, and 9.0 % for the number of unserved customers.

Experiments with 6 *vehicles*

Again, the new strategy leads to larger improvements when compared to scenario

1. While the two algorithms both succeed in servicing all customers, the new strategy reduces the objective value in nine cases. The reduction ranges from 2.0 % to 11.0 % with an average of 3.0 %.

In conclusion, the experiments show that the proposed waiting strategy is effective even if it is incorporated into a powerful algorithm, and even if it is enforced only a few times over the scheduling horizon, namely, one to four times per problem for scenario 1, and one to six times per problem for scenario 2. This strategy is also more effective when it is applied on harder problems (i.e., smaller fleet size or higher demand rates).

$\alpha_1 = 0.20$	$\alpha_1 = 0.20$	$\alpha_1 = 0.20$	$\alpha_1 = 0.20$
$s_1 = 0.60$	$s_1 = 0.80$	$s_1 = 0.80$	$s_1 = 0.80$
$\eta_1 = 200$	$\eta_1 = 100$	$\eta_1 = 200$	$\eta_1 = 300$
$V_1 = 1$	$V_1 = 1$	$V_1 = 1$	$V_1 = 1$
$\Delta_1 = 300$	$\Delta_1 = 180$	$\Delta_1 = 300$	$\Delta_1 = 600$
4.00 [a]	4.66	**4.00**	4.66
227.813 [b]	224.70	**229.550**	223.936
234.413 [c]	244.67	**179.980**	207.950
462.226 [d]	469.46	**409.530**	431.890

$\alpha_1 = 0.20$	$\alpha_1 = 0.20$	$\alpha_1 = 0.30$	$\alpha_1 = 0.40$
$s_1 = 0.90$	$s_1 = 0.90$	$s_1 = 0.90$	$s_1 = 0.80$
$\eta_1 = 200$	$\eta_1 = 300$	$\eta_1 = 200$	$\eta_1 = 100$
$V_1 = 1$	$V_1 = 2$	$V_1 = 2$	$V_1 = 1$
$\Delta_1 = 300$	$\Delta_1 = 600$	$\Delta_1 = 300$	$\Delta_1 = 180$
4.00 [a]	4.66	4.33	4.66
225.810 [b]	225.353	229.623	235.08
223.780 [c]	271.556	244.913	248.92
449.590 [d]	496.909	474.536	484.01

$\alpha_1 = 0.50$	$\alpha_1 = 0.50$	$\alpha_1 = 1.00$	$\alpha_1 = 1.00$	$\alpha_1 = 1.00$
$s_1 = 0.60$	$s_1 = 0.90$	$s_1 = 0.80$	$s_1 = 0.80$	$s_1 = 0.90$
$\eta_1 = 200$	$\eta_1 = 200$	$\eta_1 = 200$	$\eta_1 = 200$	$\eta_1 = 200$
$V_1 = 1$	$V_1 = 1$	$V_1 = 1$	$V_1 = 2$	$V_1 = 1$
$\Delta_1 = 300$	$\Delta_1 = 300$	$\Delta_1 = 300$	$\Delta_1 = 300$	$\Delta_1 = 300$
4.33 [a]	4.66	4.66	4.33	4.33
229.588 [b]	226.335	230.051	223.710	222.15
204.350 [c]	209.080	173.428	223.000	218.16
433.940 [d]	435.415	403.480	446.710	440.310

[a]Nb. of unserved customers.
[b]Total travel time.
[c]Total lateness.
[d]Objective value.

Table 3: *Best parameter values with 4 vehicles for scenario 1*

$\alpha_2 = 0.20$	$\alpha_2 = 0.20$	$\alpha_2 = 0.20$	$\alpha_2 = 0.20$	$\alpha_2 = 0.20$
$s_2 = 0.65$	$s_2 = 0.65$	$s_2 = 0.80$	$s_2 = 0.80$	$s_2 = 0.80$
$\eta_2 = 60$	$\eta_2 = 200$	$\eta_2 = 100$	$\eta_2 = 180$	$\eta_2 = 200$
$V_2 = 1$	$V_2 = 1$	$V_2 = 1$	$V_2 = 1$	$V_2 = 2$
$\Delta_2 = 100$	$\Delta_2 = 300$	$\Delta_2 = 300$	$\Delta_2 = 300$	$\Delta_2 = 300$
21.33 [a]	22.66	22.66	22	22.66
210.30 [b]	211.48	217.03	214.01	213.08
2423.00 [c]	2385.84	2357.01	2344.13	2420.80
2633.20 [d]	2597.34	2574.10	2558.02	2633.89

$\alpha_2 = 0.40$	$\alpha_2 = 0.40$	$\alpha_2 = 0.40$	$\alpha_2 = 0.40$	$\alpha_2 = 0.40$
$s_2 = 0.60$	$s_2 = 0.80$	$s_2 = 0.80$	$s_2 = 0.80$	$s_2 = 0.80$
$\eta_2 = 60$	$\eta_2 = 60$	$\eta_2 = 60$	$\eta_2 = 100$	$\eta_2 = 100$
$V_2 = 1$	$V_2 = 1$	$V_2 = 1$	$V_2 = 1$	$V_2 = 1$
$\Delta_2 = 100$	$\Delta_2 = 100$	$\Delta_2 = 180$	$\Delta_2 = 180$	$\Delta_2 = 300$
21.66 [a]	21.66	21.66	**21**	22
217.00 [b]	212.20	216.72	**217.10**	214.64
2345.50 [c]	2346.00	2489.40	**2345.90**	2371.58
2562.40 [d]	2558.14	2706.01	**2563.13**	2586.22

$\alpha_2 = 0.40$	$\alpha_2 = 0.40$	$\alpha_2 = 0.40$	$\alpha_2 = 0.40$	$\alpha_2 = 0.40$
$s_2 = 0.80$	$s_2 = 0.80$	$s_2 = 0.80$	$s_2 = 0.85$	$s_2 = 0.90$
$\eta_2 = 180$	$\eta_2 = 200$	$\eta_2 = 200$	$\eta_2 = 60$	$\eta_2 = 100$
$V_2 = 1$	$V_2 = 1$	$V_2 = 2$	$V_2 = 1$	$V_2 = 1$
$\Delta_2 = 300$	$\Delta_2 = 300$	$\Delta_2 = 300$	$\Delta_2 = 100$	$\Delta_2 = 180$
22.66 [a]	22.66	22.66	21.33	22
217.30 [b]	215.42	215.24	216.50	213.30
2429.01 [c]	2424.14	2397.37	2357.50	2573.00
2646.10 [d]	2639.57	2612.62	2574.05	2786.30

[a]Nb. of unserved customers.
[b]Total travel time.
[c]Total lateness.
[d]Objective value.

Table 4: *Best parameter values with 4 vehicles for scenario 2*

Instance number	*Number of customers*	*Number of unserved customers*	*Travel Time*	*Lateness*	*Objective Value*
1	152	12	243.24	167.67	410.91
2	140	1	244.58	229.06	473.65
3	127	0	200.84	143.21	344.05
4	144	0	240.81	281.25	522.06
5	132	0	230.15	119.24	349.39
6	120	0	207.83	221.58	429.41
7	126	0	233.27	57.97	291.24
8	155	3	243.96	801.50	1045.47
9	153	3	243.99	436.17	680.16
10	132	0	232.95	89.88	322.84
Average	138.10	1.9	232.16	254.75	486.92

Table 5: *Original algorithm under scenario 1: Simulation of 5 hours with 4 vehicles and 27 requests per hour on average.*

Instance number	*Number of customers*	*Number of unserved customers*	*Travel Time*	*Lateness*	*Objective Value*
1	152	11	233.59	196.43	430.09
2	140	1	243.74	230.13	473.87
3	127	0	202.67	137.67	340.35
4	144	0	227.30	247.09	474.40
5	132	0	219.55	120.24	339.69
6	120	0	201.78	231.42	433.20
7	126	0	224.81	52.09	276.91
8	155	2	239.43	753.19	992.62
9	153	1	240.78	459.54	700.20
10	132	0	238.67	44.49	283.17
Average	138.10	1.5	227.23	247.93	474.48

Table 6: *New algorithm under scenario 1: Simulation of 5 hours with 4 vehicles and 27 requests per hour on average.*

Instance number	*Number of customers*	*Number of unserved customers*	*Travel Time*	*Lateness*	*Objective Value*
1	152	3	248.11	25.84	273.98
2	140	0	232.64	85.79	318.43
3	127	0	218.84	90.70	309.50
4	144	0	233.54	93.37	327.91
5	132	0	226.65	29.47	256.13
6	120	0	198.02	107.32	305.34
7	126	0	217.74	29.93	247.67
8	155	0	255.32	52.53	307.85
9	153	0	235.88	40.82	276.71
10	132	0	221.48	21.69	243.18
Average	138.10	0.3	228.83	57.85	286.60

Table 7: *Original algorithm under scenario 1: Simulation of 5 hours with 6 vehicles and 27 requests per hour on average.*

Instance number	*Number of customers*	*Number of unserved customers*	*Travel Time*	*Lateness*	*Objective Value*
1	152	0	240.24	54.57	294.82
2	140	0	236.13	83.99	320.12
3	127	0	219.57	67.48	287.06
4	144	0	224.68	91.90	316.58
5	132	0	208.29	40.21	248.51
6	120	0	198.49	103.53	302.02
7	126	0	199.25	28.88	228.14
8	155	0	243.26	42.44	285.71
9	153	0	223.01	48.26	271.30
10	132	0	220.61	24.84	245.45
Average	138.10	0	221.35	58.61	279.97

Table 8: *New algorithm under scenario 1: Simulation of 5 hours with 6 vehicles and 27 requests per hour on average.*

Instance number	*Number of customers*	*Number of unserved customers*	*Travel Time*	*Lateness*	*Objective Value*
1	168	7	236.80	1258.40	1495.13
2	196	25	231.80	3077.60	3309.30
3	197	36	192.25	3249.91	3442.20
4	177	17	231.50	2158.10	2389.53
5	182	21	207.00	2732.53	2939.50
6	198	32	214.10	3395.54	3609.62
7	195	38	180.34	3284.86	3465.20
8	164	7	223.92	1391.63	1615.55
9	187	29	188.70	3275.00	3464.00
10	193	29	209.65	2374.30	2584.00
Average	187.00	24.10	211.60	2620.00	2831.40

Table 9: *Original algorithm under scenario 2: Simulation of 5 hours with 4 vehicles and 36 requests per hour on average.*

Instance number	*Number of customers*	*Number of unserved customers*	*Travel Time*	*Lateness*	*Objective Value*
1	168	7	241.40	1201.76	1443.24
2	196	21	218.72	2578.63	2797.35
3	197	35	191.40	3257.40	3448.82
4	177	14	206.70	2119.40	2326.11
5	182	21	203.01	2518.18	2721.20
6	198	30	204.55	3419.40	3623.96
7	195	38	181.90	3178.28	3360.19
8	164	6	232.00	1277.83	1509.85
9	187	27	192.85	3206.80	3399.66
10	193	28	205.63	2286.77	2492.41
Average	187.00	22.07	207.82	2504.44	2712.28

Table 10: *New algorithm under scenario 2: Simulation of 5 hours with 4 vehicles and 36 requests per hour on average.*

Instance number	*Number of customers*	*Number of unserved customers*	*Travel Time*	*Lateness*	*Objective Value*
1	168	0	251.00	283.55	534.55
2	196	0	301.10	577.42	878.50
3	197	0	274.93	1194.33	1469.30
4	177	0	262.10	567.40	829.43
5	182	0	270.54	440.02	710.60
6	198	0	265.32	560.43	825.80
7	195	0	295.90	679.10	974.96
8	164	0	256.40	364.43	621.00
9	187	0	273.80	665.00	940.00
10	193	0	246.30	557.01	803.30
Average	187.00	0	270.00	589.01	859.04

Table 11: *Original algorithm under scenario 2: Simulation of 5 hours with 6 vehicles and 36 requests per hour on average.*

Instance number	*Number of customers*	*Number of unserved customers*	*Travel Time*	*Lateness*	*Objective Value*
1	168	0	265.00	259.21	524.02
2	196	0	280.78	570.55	851.33
3	197	0	269.19	1377.41	1646.61
4	177	0	253.80	519.20	773.01
5	182	0	282.27	371.80	654.07
6	198	0	274.25	521.44	795.70
7	195	0	294.88	650.37	945.25
8	164	0	239.67	345.90	585.57
9	187	0	265.76	601.46	867.23
10	193	0	245.12	470.71	715.84
Average	187.00	0	267.07	568.80	835.86

Table 12: *New algorithm under scenario 2: Simulation of 5 hours with 6 vehicles and 36 requests per hour on average.*

6 Conclusion

In this paper, a new strategy which exploits probabilistic knowledge about future demand arrivals to better manage the fleet of vehicles is proposed. Some issues related to this strategy were addressed and investigated. Then, experiments were performed to assess its effectiveness within a previously reported tabu search heuristic. The results show that the proposed strategy provides significant improvements over the original algorithm, especially in some "critical" situations (i.e., small fleet size and high demand rates).

Acknowledgments

This work was partially supported by the Canadian Natural Sciences and Engineering Research Council (NSERC) and by the Québec Fonds pour la Formation de Chercheurs et l'Aide à la Recherche. This support is gratefully acknowledged.

References

[1] **D. Bertsekas and J. Tsitsiklis,** *"Neuro-Dynamic Programming"*, Athena Scientific, Belmont, MA. (1996).

[2] **D.J. Bertsimas,** *"Probabilistic Combinatorial Optimization Problems"*, Ph.D. Dissertation, Report No. 193, Operations Research Center, Massachusetts Institute of Technology, 1988.

[3] **D.J. Bertsimas,** *"A Vehicle Routing Problem with Stochastic Demand"*, Operations Research **40**, 574-585 (1992).

[4] **D.J. Bertsimas and L.H. Howell,** *"Further Results on the Probabilistic Traveling Salesman Problem"*, Eur. J. Oper. Res. **65**, 68-95 (1993).

[5] **D.J. Bertsimas, P. Jaillet and A.R. Odoni,** *"A Priori Optimization"*, Operations Research **38 (6)**, 1019-1033 (1990).

[6] **D.J. Bertsimas and D. Simchi-Levi,** *"A New Generation of Vehicle Routing Research: Robust Algorithms, Addressing Uncertainty"*, Operations Research **44**, 286-304 (1996).

[7] **D. Bertsimas and G. Van Ryzin,** *A Stochastic and Dynamic Vehicle Routing Problem in the Euclidian Plane*, Operations Research **39**, 601-615 (1991).

[8] **D.J. Bertsimas and G. Van Ryzin,** *Stochastic and Dynamic Vehicle Routing in the Euclidian Plane with Multiple Capacitated vehicles*, Operations Research **41**, 60-76 (1993).

[9] **M. Dror, G. Laporte and P. Trudeau,** *"Vehicle Routing with Stochastic Demands: Properties and Solution Frameworks"*, Transportation Science **23 (3)**, 166-176 (1989).

[10] **M. Gendreau, F. Guertin, J.Y. Potvin and É. Taillard,** *"Parallel Tabu Search for Real-Time Vehicle Routing and Dispatching"*, Transportation Science **33 (4)**, 381-390 (1999).

[11] **M. Gendreau, G. Laporte, and R. Séguin,** *"A Tabu Search Heuristic for the Vehicle Routing Problem with Stochastic Demands and Customers"*, Operations Research **44**, 469-477 (1996).

[12] **M. Gendreau, G. Laporte, and R. Séguin,** *"An exact Algorithm for the Vehicle Routing Problem with Stochastic Customers and Demands"*, Transportation Science **29**, 143-155 (1995).

[13] **M. Gendreau, G. Laporte, and R. Séguin,** *"Stochastic Vehicle Routing"*, Eur. J. Oper. Res. **88**, 3-12 (1996).

[14] **F. Glover and M. Laguna,** *Tabu Search*, Kluwer, 1997.

[15] **P. Jaillet,** *"Probabilistic Traveling Salesman Problems"*, Ph.D. Dissertation, Massachusetts Institute of Technology, 1985.

[16] **P. Jaillet,** *"A Priori Solution of a Traveling Salesman Problem in which a Random Subset of the Customers Are Visited"*, Operations Research **36**, 929-936 (1988).

[17] **A. Jezequel,** *"Probabilistic Vehicle Routing Problems"*, Master's Thesis, Massachusetts Institute of Technology, 1985.

[18] **A.J. Kleywegt,** *"Dynamic and Stochastic Models with Freight Distribution Applications"*, Ph.D. Dissertation, Massachusetts Institute of Technology, 1996.

[19] **G. Laporte, F.V. Louveaux and H. Mercure,** *"A Priori Optimization of the Probabilistic Traveling Salesman Problem"*, Operations Research **42**, 543-549 (1994).

[20] **A. Larsen,** *"The Dynamic Vehicle Routing Problem"*, Ph.D. Dissertation, Technical University of Denmark, 2000.

[21] **W. Powell,** *A Comparative Review of Alternative Algorithms for the Dynamic Vehicle Allocation Problem*, in Vehicle Routing: Methods and Studies,

B.L. Golden and A.A. Assad (eds), pp. 249-291, North Holland: Amsterdam (1988).

[22] **W.B. Powell,** *"A Stochastic Formulation of the Dynamic Assignment Problem, with an Application to Truckload Motor Carriers"*, Transportation Science **30 (3)**, 195-219 (1996).

[23] **W.B. Powell, G. Godfrey, K. Papadaki, M. Spivey and H. Topaloglu,** *"Adaptive Dynamic Programming for Multistage Stochastic Resource Allocation"*, Department of Operations Research and Financial Engineering, Princeton University (2000).

[24] **W.B. Powell, P. Jaillet and A. Odoni,** *Stochastic and Dynamic Networks and Routing*, in Network Routing, Handbooks in Operations Research and Management Science, Volume 8, M.O. Ball, T.L. Magnanti, C.L. Monma and G.L. Nemhauser (eds), pp. 141-295, North-Holland: Amsterdam (1995).

[25] **Y. Rochat and E. Taillard,** *"Probabilistic Diversification and intensification in Local Search for Vehicle Routing"*, Journal of Heuristics, **1**, 147-167 (1995).

[26] **R. Rubio Maqueo,** *"Dynamic-Stochastic Vehicle Routing and Inventory Problem"*, Ph.D. Dissertation, Massachusetts Institute of Technology, 1995.

[27] **R. Séguin,** *"Problèmes Stochastiques de Tournées de véhicules"*, Ph.D. Dissertation, Département d'Informatique et de Recherche Opérationnelle, Université de Montréal (1994).

[28] **M.R. Swihart and J.D. Papastavrou,** *"A Stochastic and Dynamic Model for the Single-Vehicle Pick-up and Delivery Problem"*, Eur. J. Oper. Res. **114**, 447-464 (1999).

[29] **É. Taillard, P. Badeau, M. Gendreau, F. Guertin and J. Y. Potvin,** *"A Tabu Search Heuristic for the Vehicle Routing Problem with Soft Time Windows"*, Transportation Science **31**, 170-186 (1997).

[30] **W.H. Yang, K. Mathur and R.H. Ballou,** *"Stochastic Vehicle Routing Problem with Restocking"*, Transportation Science **34 (1)**, 99-112 (2000).

[31] **C.D.J. Waters,** *"Vehicle Routing Problems with Uncertainty and Omitted Customers"*, Journal of the Operational Research Society **40**, 1099-1108 (1989).